

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: RENDERING A SILHOUETTE EDGE
APPLICANT: ADAM T. LAKE, CARL S. MARSHALL AND MICHAEL
MESNIER

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL 935340463US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the U.S. Patent and Trademark Office, P.O. Box 2327, Arlington, VA 22202.

Date of Deposit

11-28-91

Signature

Gabe Lewis

Typed or Printed Name of Person Signing Certificate

RENDERING A SILHOUETTE EDGE

Background of the Invention

5 This invention relates to rendering a silhouette edge of a three-dimensional ("3D") model.

10 Silhouette edges are outlines of a 3D model and its interior features. Silhouette edges are detected based on the intersection of front and back facing polygons that make up the 3D model. Computer programs are used to detect, and render, silhouette edges to form a two-dimensional ("2D") image from the 3D model. How those edges are rendered affects the look of the 2D image.

Summary of the Invention

15 In general, in one aspect of the invention, silhouettes edges of a 3D model are detected and rendered in a format that corresponds to the geometry of the 3D model.

20 Among the advantages of the invention are one or more of the following. Rendering the silhouette edges in a format that corresponds to the geometry of the 3D model produces a 2D image that more closely approximates the 3D model. For example, curvature in the original 3D model can be detected, and curved silhouette edges can be selected to preserve the curvature of the original 3D model.

Other advantages and features will become apparent from the following description and the claims.

Brief Description of the Drawings

Fig. 1 is a view of a 3D model.

5 Fig. 2 is a view of polygons in the 3D model.

Fig. 3 is a view of one of the polygons.

Fig. 4 is a flowchart showing a process for rendering a 2D image using silhouette edges in accordance with one embodiment of the invention.

10 Fig. 5 is a perspective view of polygons which illustrates how silhouette edges are detected.

Fig. 6 is a view of adjacent polygons.

Fig. 7 is a view of silhouette edges.

15 Figs. 8a, 8b and 8c are texture maps for rendering the silhouette edges.

Fig. 9 is a flowchart showing a texture mapping process in accordance with one embodiment of the invention.

Fig. 10 shows a texture map area.

20 Fig. 11a shows a 2D image rendered with all straight silhouette edges, and Fig. 11b shows the same 2D image rendered with texture-mapped silhouette edges.

Fig. 12 is a view of a computer system on which the processes of Figs. 4 and 9 may be implemented.

Description

As shown in Figs. 1 and 2, 3D model 10 is comprised of interconnecting polygons 11. Polygons 11 are triangles in this embodiment; however, other types of polygons may be used. Groups of polygons are organized into meshes, each of which corresponds to an element of the 3D model. For example, arm band 14 and belt 15 are in different meshes. In a mesh, adjacent polygons share common features. For example, in arm band mesh 16, polygons 17 and 19 share a common edge 20 (see close-up view 23).

The 3D data for each triangular polygon is comprised of three vertices positioned in (x,y,z) coordinate space. These vertices define a face for each polygon in the image. The 3D data also includes a normal vector to each face, which is also positioned in (x,y,z) coordinate space. For example, in Fig. 3, polygon 24 includes face 25, vertices 26 to 28, and normal vector ("N") 29. This 3D data is used in process 30 of Fig. 4 for rendering non-photorealistic ("NPR") images (such as cartoon-like images) of 3D model 10 using silhouette edges.

Process 30 begins in 400 by creating an edge list from the 3D data. To do this, in 402, the face of a polygon in the 3D model is examined and the edge coordinates of the polygon are detected based on its vertices. In 403, the

extracted edge coordinates are stored in an edge list in a hash table. The hash table includes entries that are accessible via keys stored in correspondence with the entries. Vertices of the edges are used to determine the keys for the hash table. The keys are determined by adding coordinates for the vertices of an edge, and dividing the resulting sum by the total size of the hash table. The remainder (or "Mod") of the resulting quotient is used as the key to access each edge in the hash table.

10 In 401, process 30 detects silhouette edges in the 3D data based on the edge list. In particular, process 30 determines, in 404, which edges in the hash table are silhouette edges. Two types of silhouette edges are detected by process 30. One type of silhouette edge is a common edge shared by two polygon faces -- a front face and a back face. This type of silhouette edge defines a transition point between front and back features of an image relative to a preset point (the "eyepoint") on a display screen. An eyepoint vector is used to determine whether a polygon face is front-facing or back facing. The eyepoint vector is the vector from the preset point to the face of a target polygon. An eyepoint vector 36 (pointing out of the page) and a normal vector 37 are shown in Fig. 1.

To determine whether an edge 39 shared by two polygons 40 and 41 (Fig. 5) is a silhouette edge, process 30 executes the steps of block 42. In 405, process 30 determines the vector dot product of normal vector 44 (to polygon face 45) and eyepoint vector 46. In 406, process 30 determines the vector dot product of normal vector 47 (to polygon face 49) and eyepoint vector 46. The signs of the resulting dot products are then compared in 407. If the signs are the same, this means that each of polygon faces 45 and 49 can be viewed from the eyepoint on the screen (or they are both back-facing). If this were the case, edge 39 would not be a silhouette edge. If the signs are different, the two polygon faces 45 and 49 are facing in different directions, meaning that one of the polygon faces cannot be viewed from the eyepoint. This is the case in Fig. 5; therefore, edge 39 is a silhouette edge.

As shown in Fig. 6, edges 150, 151, 152, 153, 154 and 155 of image 48 are silhouette edges since they are shared with polygons that cannot be seen from the eyepoint (for example, 90° to the page showing Fig. 6). Edges 156, 157, 158, 159 and 160 are not silhouette edges since they are shared by polygons that can be seen from the eyepoint.

Another type of silhouette edge is found on a boundary of a mesh. This type of silhouette edge is not

shared with an adjacent polygon. For example, in Figs. 1 and 2, polygons on the perimeter of arm band 14 (such as polygon 21) contain at least one edge (edge 22) that is not shared with an adjacent polygon in the same mesh (though such an edge may abut a polygon of another mesh). To detect this type of silhouette edge, in 408, process 30 detects edges in a mesh that are not shared with an adjacent polygon in the same mesh. Detecting can be performed by comparing vertices of polygon edges to determine if the same vertices are used in two different polygons.

Once a silhouette edge has been detected, in 409 it is marked in the hash table containing the list of all edges. This may be done after each silhouette edge is detected, or all silhouette edges in an image may be detected and then marked. The silhouette edges are marked with a boolean flag; however, other markings may be used. In 409a, process 30 renders the 3D model.

In 410, process 30 renders an NPR 2D image from the silhouette edges marked in the hash table. The marked silhouette edges are retrieved from the hash table in 411 and rendered in 412 to 414 in a format that corresponds to the geometry of the original 3D model. What is meant by "geometry" is the curvature, texture, contours or other physical features of the 3D model. Rather than rendering

all individual edges (or edge portions) as straight lines, the geometry of the 3D model is used to determine the display format of the individual edges. For example, if the geometry of the 3D model includes a curve, then a curved line is used in the silhouette edge rendering process.

One way of determining the geometry of the 3D model is to determine the angle between adjacent silhouette edges. In 412, process 30 detects this angle based on a silhouette edge 51 and its adjacent silhouette edge 54 (Fig. 7).

Process 30 determines the dot product of edges 51 and 54 in 416. This dot product is the cosine of angle 55 between silhouette edges 51 and 54. To determine this angle, in 417 process 30 determines the inverse cosine of the dot product.

Once the angle between adjacent silhouette edges has been determined, process 30 selects (in 413) a display format for the target silhouette edge based on the angle. In this embodiment, the display format is a texture map; however, other display formats may be used. In a simple example, one of the texture maps shown in Figs 8a, 8b and 8c may be selected. Fig. 8a shows a straight-stroke texture map 56 for generating straight silhouette edges. Figs. 8b and 8c show respective leftward-stroke 57 and rightward-stroke 59 texture maps for generating curved silhouette

edges. Which of these texture maps is selected by process 30 in 413 depends upon the angle determined in 412.

Straight-stroke texture map 56 is selected in 413 if an absolute value of the angle determined in 412 is less than a predetermined value, such as 10° . Other values, however, may be used. If the absolute value of the angle is greater than 10° and the angle is positive (for example), leftward-stroke texture map 57 is selected in 413. If the absolute value of the angle is greater than 10° and the angle is negative (for example), rightward-stroke texture map 59 is selected in 413.

In 414, process 30 displays the silhouette edge using the selected texture map. Fig. 9 shows a process 60 for rendering the silhouette edge using texture mapping. In 901, process 60 determines a size of an area into which the texture map is to be rendered (the "texture map area"). The texture map area is a quadrilateral, such as a rectangle, having a width that is orthogonal to its height and to the eyepoint vector. The texture map area is configured this way so that all texture maps face the eyepoint and thus can be viewed from the same perspective.

To determine the size of the texture map area, in 902 process 60 constructs a normal vector relative to a silhouette edge, in this example, between two end points 62

and 64 of silhouette edge 65 (see Fig. 10). Normal vector 66 defines the height of the texture map area. In 903, process 60 determines the cross product of normal vector 66 and an eyepoint vector 67. This cross product produces a
5 width vector 69, which defines the direction of the width of the texture map area. The width vector is orthogonal to both normal (height) vector 66 and eyepoint vector 67.

In 904, process 60 defines the texture map area based on four map points that are positioned relative to end
10 points 62 and 64 in a direction of width vector 69. The four map points 70a to 70d are positioned so that one map point is on each side of, and equidistant from, each end point (see Fig. 10). The distance between each map point and a corresponding end point determines the size of texture
15 map area 71. This size may be pre-set, or it may be controlled by a computer animator via a graphical user interface (not shown). The size is not necessarily dictated by the magnitude of width vector 69.

Once the texture map area is defined for each
20 silhouette edge, the texture map selected for that area in 413 (of process 30) is applied in 905. Display of the silhouette edges using texture mapping may be performed for each silhouette edge immediately after a texture map is selected for that edge. Alternatively, texture maps may be

selected for all of the silhouette edges of an image (or some subset of silhouette edges, such as a mesh), and then texture mapping may be performed.

Figs. 11a and 11b show how processes 30 and 60 affect the look of a 2D NPR image. Image 74 in Fig. 11a is an NPR image that was rendered using all straight lines for the silhouette edges. Image 75 in Fig. 11b is an NPR image that was rendered using processes 30 and 60 and texture maps 56, 57 and 59 (Figs. 8a to 8c). Comparing the two images 74 and 75, it is clear that the lines in image 75 are relatively fluid, and resemble a cartoonist's sketch more so than those in image 74.

Processes 30 and 60 may be used to create animation cels for cartooning. For example, a 3D model, such as model 10, may be generated, and then positioned in a desired manner. Processes 30 and 60 may be executed on the model to produce an NPR 2D image for that position. Then, the 3D model 10 can be re-positioned (e.g., rotated), and processes 30 and 60 executed on the re-positioned model to produce an NPR 2D image for the new position. This process may be repeated to produce NPR 2D images for any number of positions. In this way, processes 30 and 60 can generate animation cels automatically.

Processes 30 and 60 are performed in 3D space, meaning that they are performed using 3D data. By processing data in 3D space (as opposed to 2D space), processes 30 and 60 can be performed using a processor and rendering can be done using graphics accelerator cards of a personal computer.

Fig. 12 shows a computer 80 for rendering NPR images using processes 30 and 60. Computer 80 includes a processor 81, a memory 82, a storage medium 84 (e.g., a hard disk), and a 3D graphics accelerator card 85 for repositioning a 3D model and manipulating 3D data (see view 86). Storage medium 84 stores hash table 87 containing the edge list described above, 3D data 89, and computer instructions 90 for rendering NPR 2D images via processes 30 and 60.

Processes 30 and 60, however, are not limited to use with any particular hardware or software configuration; they may find applicability in any computing or processing environment. Processes 30 and 60 may be implemented in hardware, software, or a combination of the two. Processes 30 and 60 may be implemented in computer programs executing on programmable computers that each include a processor, a storage medium readable by the processor (including volatile and non-volatile memory and/or storage elements), at least one input device, and one or more output devices. Program

code may be applied to data entered using an input device to perform processes 30 and 60 and to generate output information. The output information may be applied to one or more output devices, such as display screen 91.

5 Each such program may be implemented in a high level procedural or object-oriented programming language to communicate with a computer system. However, the programs can be implemented in assembly or machine language. The language may be a compiled or an interpreted language.

10 Each computer program may be stored on a storage medium or device (e.g., CD-ROM, hard disk, or magnetic diskette) that is readable by a general or special purpose programmable computer for configuring and operating the computer when the storage medium or device is read by the
15 computer to perform processes 30 and 60. Processes 30 and 60 may also be implemented as a computer-readable storage medium, configured with a computer program, where, upon execution, instructions in the computer program cause the computer to operate in accordance with processes 30 and 60.

20 Other embodiments are also within the scope of the following claims. For example, methods of detecting silhouette edges other than that described above in process 30 may be used during the NPR image rendering process.

The invention is not limited to the texture maps shown in Figs. 8a, 8b and 8c. Other types (e.g., shapes) and styles of texture maps may be used. For example, the texture maps shown in Figs. 8a, 8b and 8c may be implemented
5 in brush-style, pencil-style, or crayon-style.

Alternatively, a single texture map may be used to render all silhouette edges of a 3D model. For example, a charcoal-style silhouette edge could be used to give the resulting image a "sketchy" feel. In cases where a single
10 texture map is used, there is no need to determine the angle between adjacent silhouette edges.

Other criteria for determining which texture maps to use to render a silhouette edge (meaning other than an angle between two adjacent silhouette edges) may also be employed,
15 depending upon the geometry of the original 3D model and the complexity of the texture maps being used.

What is claimed is: